

PAPER TWO OF THREE · ACCOUNTING AND AUDIT

Sufficient and Appropriate

Artificial intelligence, audit evidence and professional scepticism, in the first market where a regulator has said what it expects

Paul Forrest

Fractional Head of AI, Ecaveo
August 2026

About this paper

This is the second of three papers on deploying artificial intelligence inside professional firms that carry personal accountability for their work. The first concerns private equity, the third law firms. Each paper stands alone. Together they argue one thing, which is that the durable advantage in professional work comes from decisions that can be reconstructed afterwards, and that systems which produce answers without producing a record destroy more value than they create.

Audit sits in the middle of the series because it is the one market where a regulator has already said what it expects. In March 2026 the Financial Reporting Council published guidance on generative and agentic artificial intelligence, which it describes as the first issued by any global audit regulator.⁵ That changes the question facing a firm from what might be required to what has been said.

Editorial balance. Roughly seven parts in ten of this document is evidence and sector analysis, two parts in ten is original framework, and one part in ten concerns how I would run the work. Where the evidence is thin, the paper says so.

HOW TO READ THE EVIDENCE TAGS

Every substantive claim carries a tag stating what kind of evidence supports it.

REGULATOR A regulator, oversight body or public authority speaking in its own voice.

STANDARD A published auditing or management standard, or a professional requirement.

PEER-REVIEWED A primary study, working paper or systematic review.

SURVEY A survey or published dataset with disclosed method and sample.

VENDOR RESEARCH Research, a survey or a claim from a firm with a commercial interest in the answer. Directional only.

PRACTICE Practitioner opinion, including my own.

EVIDENCE GAP A question the available evidence does not answer.

ECAVEO My own framework, proposed rather than proven.

Regulatory and professional notice. Parts of this paper describe auditing standards, regulatory guidance and professional requirements, in particular in Chapters 2 and 5. Nothing in this document constitutes legal advice or an assurance opinion. The provision of legal advice sits outside the terms of any engagement with the author or with Ecaveo, and the material is presented to support discussion and further review by qualified advisers.

Statements about what a source says are attributed to that source. Where I draw a practical implication the source does not itself state, the text says so. The standards described here are still moving.

Contents

A note from the author 4

Executive summary 5

ONE The evidence question 7

TWO What the regulators have actually said 10

THREE Where the value sits across a firm 15

FOUR Six ways this fails in an audit 18

FIVE Certification, and the control model 21

SIX Proving it, and the evidence gate 24

SEVEN The economics of verification 26

EIGHT Ninety days, then nine months 28

NINE The session to run with the board 29

TEN Where the evidence is thin 30

References 32

THE ARGUMENT IN ONE PARAGRAPH

Artificial intelligence changes what an audit team can examine and leaves untouched the question of what constitutes sufficient appropriate evidence. The firms that will do well are the ones that treat a generated output as information that still needs evaluating, that certify tools before use and monitor them afterwards, and that can show an inspector the basis on which a tool was trusted. The Financial Reporting Council has already said that the human auditor is always accountable.⁵ The work now is making that statement operational.

ILLUSTRATIVE MATERIAL

Every firm, client, engagement and person in the worked examples is invented. Each is marked as an illustrative scenario where it appears, and none describes a real firm or audited entity. Published research is cited and numbered. Invented material is never presented as evidence.

CITATION

Forrest, P. (2026). *Sufficient and Appropriate: Artificial Intelligence, Audit Evidence and Professional Scepticism*. Ecaveo Responsible AI Series, Paper Two. Version 1.0, August 2026.

ALSO IN THIS SERIES

Paper one. *The Traceable Thesis*. Artificial intelligence across the private equity investment lifecycle.

Paper three. *Verified Authority*. Artificial intelligence, privilege and the duty to the court in law firms.

A NOTE FROM THE AUTHOR

The oldest question in the profession

Auditing has spent a century refining what counts as evidence. The current wave of technology has arrived with a great deal to say about production and very little to say about that.

I came to this work through business process reengineering in the 1990s and moved into machine learning and language models from 2014. I build small, specific models in preference to large general ones, on the view that a model trained on everything is calibrated for nothing in particular, and that audit work rewards narrow reliability over broad fluency.

What strikes me about the current moment in audit is how well the profession is placed and how rarely it says so. Every other sector is inventing governance vocabulary from nothing. Audit already has a quality management standard that requires appropriate technological resources, an engagement standard that puts a named partner in charge of them, and a regulator that inspects both.^{23,24} **STANDARD** The apparatus exists. What has been missing is a settled view of how a tool earns the right to be relied on, and in March 2026 the Financial Reporting Council supplied one.

There is a harder problem underneath, and it is the reason for the title. Sufficient appropriate evidence is a judgement about the relationship between what you have and what you are asserting. A system that produces a plausible summary of 400 board minutes has produced information. Whether that information is evidence depends on what was done to it, what it was checked against, and whether anyone can now demonstrate either. Firms that blur those two things will pass their own internal reviews and fail an inspection.

This paper argues that the profession's advantage lies in a discipline it already owns, which is the habit of asking what an assertion rests on. Applied to its own tools, that habit is the whole of the answer.

TWO PAGES THAT CIRCULATE ON THEIR OWN

Executive summary

A regulator has set the expectation, the standards are being rewritten around it, and the practical problem has moved from permission to proof.

The position in audit is further advanced than in any adjacent profession, and that is an advantage for firms that use it. The obligations are legible, the inspection regime is established, and the direction of the international standards is visible in draft. What remains is the work of showing that a given tool, used for a given purpose, produced something a firm was entitled to rely on.

1

The regulator has spoken, and it has not moved the accountability. The Financial Reporting Council's guidance of 30 March 2026 states that the human auditor is always accountable and that the guidance does not alter that position, and it ties responsibility to the existing quality management architecture.⁵ **REGULATOR** It also declines to prescribe a single validation protocol, treating the extent of mitigation as a matter of professional judgement that varies with the tool and its intended use.

2

Certification practice at the largest firms was examined in June 2025 and found uneven. A thematic review of the six largest UK audit firms reported that certification maturity varied, that firms were still planning artificial-intelligence-specific processes and had yet to operate them, that no firm monitored what these tools did to audit quality after deployment, and that recertification was often triggered informally by a tool owner.⁷ **REGULATOR**

3

The evidence standard itself is being rewritten, and the drafting is technology-neutral by design. The international board released exposure drafts in August 2026 covering audit evidence, risk response and analytical procedures, proposing an input and output model of evidence, strengthened evaluation of authenticity, and the replacement of the term automated tools and techniques with technological tools.⁸ **STANDARD** Comments close in December 2026, so nothing is final.

4

In the United States the requirement is already in force for current-year audits. Amendments on technology-assisted analysis of information in electronic form took effect for fiscal years beginning on or after 15 December 2025, and require the auditor to establish the reliability of company-provided electronic information and to investigate the items the analysis flags.⁹ **STANDARD** A firm with US-listed clients is inside that regime now.

5

No enforcement case yet turns on misuse of these tools. I could find no published enforcement decision in which the misuse of an artificial intelligence or automated tool was the stated ground of sanction. **EVIDENCE GAP** My own view is that the first such finding will concern documentation of the basis for reliance. A wrong answer will come later.

6

The threat to evidence is arriving from outside the firm as well as inside it. Documents, confirmations and even video calls can now be fabricated convincingly, as the reported theft of around \$25m from an engineering group through a deepfaked video conference illustrates.¹⁶ **REGULATOR** Authenticity is becoming a live audit procedure. It used to be an assumption.

What follows for a firm

Three things, in order. Separate the tools that support work from the tools that contribute to a conclusion, because the assurance burden differs by an order of magnitude and most firms currently apply one policy to both. Build the certification record the thematic review found missing, meaning what the tool does, why it suits the intended use, how it was tested, who approved it, when it will be reviewed, and what changed since.^{6,7} Then monitor after deployment, which is the step the review found firms were least able to demonstrate.

Underneath all of it sits a single discipline. An audit file should show what a generated output was checked against, and by whom. That is the only thing an inspector can test, and it is the only thing that distinguishes reliance from assumption.

Chapter 8 sets out a 90-day proof on three bounded workflows. Chapter 10 states what would change the position taken here.

CHAPTER ONE

The evidence question

Information becomes evidence when a procedure is applied to it. That distinction is old, unglamorous and the whole of the problem.

An auditor gathers information and performs procedures on it, and what survives the procedure is evidence for an assertion. Nothing in that sequence changes because the information arrived through a model. What changes is the volume, the apparent authority of the presentation, and the ease with which the procedure can be skipped without anybody noticing.

The international board has put that distinction at the centre of its current drafting. The exposure draft on audit evidence proposes what the board describes as an input and output model, under which information becomes audit evidence only once audit procedures have been applied to it, alongside strengthened requirements to evaluate the relevance, reliability and authenticity of information significant to the audit.⁸ **STANDARD** The proposal is open for comment until December 2026, so it indicates direction and requires nothing yet.

Read against practice, that framing is useful because it names the failure precisely. A team that places a generated summary of board minutes in the file, cross-referenced to nothing, has filed information and described it as evidence. A team that places the same summary alongside the minutes it cites, with the significant matters agreed back to source, has performed a procedure. The two files look similar and differ completely.

THE CORE CLAIM

A generated summary is information. It becomes evidence when a procedure has been applied to it, and the file shows what that procedure was.

1.1 The cycle this series uses

Every paper in this series describes the same five-stage cycle, because it is the sequence a professional decision follows whether or not anybody writes it down. In audit the stages have names already, which is one reason the profession is better placed here than it believes.

THE CYCLE EVERY PAPER IN THIS SERIES EXAMINES

1 Gather Information is obtained from the entity, from third parties and from the firm's own methodology and knowledge.	2 Ground A procedure is applied, and each assertion is tied to the source it rests on.	3 Weigh A qualified professional judges whether what is held is sufficient and appropriate.	4 Decide A conclusion is reached and documented so that an experienced auditor could follow it.	5 Compare Inspection, review and hindsight read the conclusion back against the evidence.
---	--	---	---	---

FIGURE 1 Audit already names all five stages. The risk in the current wave is that stage two becomes invisible because stage one has become so productive.

Ecaveo framework, used across the series. **ECAVEO**

1.2 What the technology is actually good at

Three capabilities carry most of the value in an audit practice, and each has a different evidential status. Retrieval across a standards library and a firm's own technical guidance is the safest, because the output can be checked against an authoritative text in seconds. Extraction from documents, meaning terms from contracts, matters from minutes and figures from agreements, is next, because the source is present and the check is mechanical. Analysis over full populations, meaning journals, transactions and balances, carries the most value and the most risk, because the output is a set of items to investigate.

The regulator's own illustration is an anomaly-detection tool that flags unusual journal entries in support of fraud identification, which it describes as one coherent approach among others.⁶ **REGULATOR** That is the right posture for a regulator and an awkward one for a firm, because it leaves the design choices where they belong, which is with the firm.

1.3 What the profession's own evidence says about adoption

Adoption among accountants is broad and its depth is unclear. A survey of 2,718 chartered accountants across 48 countries, published in April 2025, reported that 85 per cent were willing to use these tools, rising to 91 per cent of those aged 18 to 25, while 53 per cent of C-suite respondents felt unprepared for the effect on their organisation.¹³ **SURVEY** Willingness at 85 per cent and preparedness at 47 per cent, in the same survey, on the same page.

A separate study of 1,816 professionals across 62 countries, fielded in the spring of 2026 and covering legal, tax, audit and risk, reported that 74 per cent used these tools at least weekly and that around a third used tools their organisation had not sanctioned.¹⁴ **VENDOR RESEARCH** The second figure matters more than the first. Unsanctioned use inside an audit practice is a confidentiality exposure and a documentation gap at the same time, and neither shows up until somebody goes looking.

A third of professionals reporting unsanctioned use is not an adoption statistic. It is a control finding.

The largest firms have announced substantial investment, with figures of \$1bn, \$1.4bn and \$2bn appearing in announcements made from 2023 onward.¹⁵ **VENDOR RESEARCH** The profession's own institute maintains a practical guide to using these tools, and has published its own testing of three platforms against audit tasks, which is a more useful starting point for a mid-tier firm than any of the investment figures.^{11,12} **REGULATOR** Those are self-disclosed commitments, and what they describe is ambition. They tell a mid-tier firm very little about what to do on Monday.



Technology-assisted analysis over a full population. The output is a set of items to investigate, and the investigation is the audit procedure.

CHAPTER TWO

What the regulators have actually said

Audit is the first profession where a regulator has set out what it expects. The expectation is proportionate, which means the firm carries the design work.

A reminder on register. This chapter summarises what published sources say and grades how firmly they say it. It does not state the regulatory position of any firm. Where I draw an implication a source does not itself state, the sentence says so.

2.1 The Financial Reporting Council guidance

On 30 March 2026 the Financial Reporting Council published guidance on generative and agentic artificial intelligence in audit, describing it as the first such guidance issued by a global audit regulator.⁵ **REGULATOR** Four features matter for a firm planning its response.

The guidance is explicit that responsibility does not move. It states that the human auditor is always accountable and that the guidance does not alter that position, and it ties the point to the existing architecture by stating that the firm and the engagement partner retain full responsibility for audit quality in accordance with the quality management and engagement standards.⁵ **REGULATOR**

It identifies three categories of risk, which are the misuse or misinterpretation of outputs, deficient output arising from flawed system design, and a methodology that fails to comply with auditing standards.⁵ **REGULATOR** Worth highlighting is that only the second is a technology problem. The first is a training and supervision problem and the third is a methodology problem, which places two thirds of the regulator's concern inside functions a firm already runs.

It declines to prescribe a uniform validation protocol, stating that the nature and extent of mitigating activities is a matter of professional judgement varying with the tool and its intended use, and it describes confidence as built through four combined layers covering system design, certification, staff education and governance, and human oversight.⁵ **REGULATOR** **PRACTICE** My own view, which goes beyond what the guidance states, is that firms will find the absence of a prescribed protocol harder to work with than a prescribed one would have been, and that the second layer is where most of them will struggle.

REGULATOR

The guidance is addressed principally to central methodology and tooling functions, which tells a firm where the work sits.

2.2 The certification problem, already documented

The Council had examined certification nine months before that guidance appeared. Its thematic review of the certification of automated tools and techniques, published in June 2025 and covering the six largest UK audit firms, found that the maturity of certification processes varied and that some lacked formal documented policies.⁷ **REGULATOR** Firms acknowledged the specific risks of artificial intelligence within these tools and were, at that point, planning supplementary certification processes and had not yet begun to operate them.

Three further findings bear directly on what a firm should build. No firm performed formal monitoring to quantify what a tool did to audit quality after deployment, and the monitoring that existed was generally restricted to the more widely used tools and performed largely to understand how many teams were using them for licensing purposes. Data-validation requirements were inconsistent, with some certification templates omitting an explicit requirement to check the completeness and accuracy of input data before use. Recertification triggers were informal, typically relying on a tool owner to raise the need.⁷ **REGULATOR**

Read together with the March 2026 guidance, those findings describe the gap a firm has to close. The guidance says confidence should be built partly through certification and partly through oversight after deployment. The review, which predates the guidance and does not use its four-layer framing, found certification processes present at every firm and monitoring of audit quality impact present at none. Which of those two is the harder problem depends on the firm, though I would expect the monitoring to be. **ECAVEO**

THE REGULATORY POSITION AT A GLANCE

SOURCE	WHAT IT SAYS	STATUS
FRC guidance, 30 March 2026	The human auditor is always accountable, and confidence is built through design, certification, education and oversight ⁵	Guidance, in force, not prescriptive
FRC illustrative example, June 2025	A worked example of anomaly detection in journals, with documentation expectations ⁶	Guidance, described as one coherent approach
FRC thematic review, June 2025	Certification practice at the six largest firms varies, and monitoring after deployment is weak ⁷	Inspection finding
ISQM 1 and ISA 220 (Revised), and their UK counterparts	A quality objective addressing appropriate technological resources, and an engagement partner responsible for resources ^{23,24}	Standards, effective since December 2022
PCAOB amendments	Reliability of electronic information must be evaluated, and flagged items investigated ⁹	Requirement, effective for periods beginning on or after 15 December 2025
IAASB exposure drafts, August 2026	An input and output model of evidence, and technological tools in place of automated tools and techniques ⁸	Draft, comments close December 2026

TABLE 1 The UK position is guidance built on existing standards. The US position is a requirement already in force. The international position is in draft.

2.3 The American requirement is live now

A UK firm auditing an entity registered with the US oversight board is inside a requirement, and guidance does not describe it. Amendments adopted in June 2024 and approved that August took effect for fiscal years beginning on or after 15 December 2025, and they amend the standards on audit evidence and on responses to risk.⁹ STANDARD

Three obligations follow. Where external information is provided by the company in electronic form, the auditor obtains an understanding of its source and of the company's process for handling it, and then either tests the information for modification or tests the relevant controls. Where one procedure serves more than one purpose, each objective is separately achieved. Items that the analysis flags as meeting specified criteria are investigated to determine whether they indicate a misstatement or a control deficiency.⁹ STANDARD Supporting staff guidance on the reliability of externally provided electronic information followed in October 2025.

My own reading, which the standard does not state in these terms, is that the third obligation is the one most likely to be missed in practice. A tool that flags 4,000 items creates an investigation obligation for 4,000 items, and the instinct to reduce that population by re-running the tool with tighter parameters changes the procedure.

2.4 What is changing internationally, and what is not

Two distinct projects are often conflated. The board's ISA 500 series project covers the standards on inventory and external confirmations, and remains at the stage of developing an exposure draft following approval of the project proposal in March 2026.¹⁰ STANDARD The revision of the audit evidence standard itself sits in a separate project on audit evidence and risk response, covering risk response, audit evidence and analytical procedures, whose exposure drafts were approved in June 2026 and released in August 2026 with comments closing on 15 December 2026.⁸ STANDARD

A firm reading only the first project page will conclude that the evidence standard is untouched. It is not. The drafting proposes the input and output model described above, strengthens the evaluation of relevance, reliability and authenticity, replaces the term automated tools and techniques with technological tools, and adds requirements on professional scepticism aimed at confirmation bias, meaning the tendency to favour corroborative evidence over contradictory evidence.⁸ STANDARD None of it is final and all of it indicates direction.

2.5 What binds underneath

The quality management standard has required since December 2022 that a firm establish a quality objective addressing appropriate technological resources, obtained or developed, implemented, maintained and used, both for the firm's own system and for the performance of engagements.²³ **STANDARD** The engagement standard requires the engagement partner to determine that sufficient and appropriate resources are assigned or made available in a timely manner.²⁴ **STANDARD**

The American standard on audit evidence has recognised since 2020 that automated tools and techniques, including artificial intelligence, are a means of obtaining audit evidence, and requires evaluation of information regardless of its source, including its susceptibility to bias.²⁰ **STANDARD** Firms outside the United States will find it a useful cross-check on their own methodology.

Outside the audit standards, the same general obligations apply as to any organisation. Data protection reaches client financial, payroll and personal information, and the Information Commissioner's guidance remains the principal reference although the Commissioner has flagged it as under review.¹ **REGULATOR** The security guidelines co-sealed by agencies from 18 countries set expectations across design, development, deployment and operation, and the American standards body's generative profile sets out 12 risk categories.^{2,3} **STANDARD** A certifiable management system exists in ISO/IEC 42001.²² **STANDARD**

A photograph of three business professionals in a modern office setting. A man with glasses and a beard, wearing a dark suit, is seated on the left, looking at a laptop screen displaying a bar chart. A woman with long dark hair, wearing a dark blazer, is standing in the center, leaning over a desk and pointing at a document with a pen. Another woman with blonde hair, wearing a dark blazer, is seated on the right, holding a document and looking at it. The desk is cluttered with papers, a laptop, a mug, and a glass. In the background, a large window offers a view of a city skyline, including a prominent domed building. The overall atmosphere is professional and collaborative.

PART TWO

Where the value sits, and how it fails

CHAPTER THREE

Where the value sits across a firm

The strongest cases sit where an authoritative answer already exists to check against. The weakest sit where the output is a conclusion.

A practice contains work of three kinds, and the distinction should drive the whole portfolio. There is work where a correct answer exists and can be consulted, such as what a standard says. There is work where a source exists and the output can be agreed back to it, such as what a contract provides. There is work where the output is a judgement, such as whether a balance is fairly stated. The first is where deployment should start and the third is where it should not go.

THE OPPORTUNITY PORTFOLIO

WORKFLOW	WHAT THE TECHNOLOGY DOES WELL	THE VALUE TEST
Technical research	Retrieves current standards, firm guidance and interpretations, with citations	Speed to a supported answer, measured on jurisdiction, effective date and citation accuracy
Evidence extraction	Pulls terms from contracts, matters from minutes and data from agreements	Recall of known items, false positives and the time a reviewer spends agreeing back to source
Population analysis	Examines journals, transactions and balances for anomalies	Coverage and the proportion of flagged items that repay investigation
Working papers	Drafts descriptions, cross-references and responses to review notes from approved evidence	Consistency and reviewer effort, with no assertion lacking an evidential basis
Client accounting	Classifies transactions, flags anomalies and drafts management commentary	Close cycle time, with human approval of every posting
Firm operations	Supports proposals, learning, policy search, scheduling and knowledge management	Low-risk productivity and better reuse of what the firm already knows

TABLE 2 The value test is stated in place of the benefit. A benefit nobody measured against a baseline will not survive its first quality review.

3.1 Technical research is the safest place to start

A technical accounting assistant restricted to approved current standards and firm guidance has the property that makes evaluation cheap, which is a checkable answer. A firm can assemble 100 questions whose answers are known, measure citation accuracy, jurisdiction accuracy and effective-date accuracy, and measure how often the system declines to answer when it should.

That last measure deserves more attention than it usually receives. A research assistant that answers everything is worse than one that answers 80 per cent and says so, because the 20 per cent it should have declined arrives with the same confident tone as the rest. My own practice is to treat the rate of appropriate abstention as a first-class acceptance criterion. **PRACTICE**

3.2 Extraction, and the discipline of agreeing back

Extraction is where most of the near-term hours sit. Revenue contract review, lease data, board minutes and loan agreements all involve reading documents for defined items. All of them are slow by hand.

The control that makes extraction safe is unglamorous. Every extracted item carries a pointer to the clause or paragraph it came from, and the reviewer agrees a sample back to source and leaves the rest. That converts review from re-performance into testing, which is the profession's own method applied to its own tools.

ILLUSTRATIVE SCENARIO

The minutes that were summarised accurately

An invented mid-tier firm uses a tool to extract audit-relevant matters from 18 months of board and committee minutes for a manufacturing client. The summary is accurate on every matter it reports.

The audit committee met five times in the period and the pack for the fourth meeting was never uploaded, because it had been circulated by a different secretary under a different naming convention. The tool reported on four meetings. Nobody established that there had been five.

The extraction was flawless. Nobody tested completeness, which was the assertion that mattered.

3.3 Population analysis, and the obligation it creates

Examining a full population is the most attractive claim in this field and the one that carries the heaviest obligation. Where a tool flags items meeting specified criteria, the US requirement is to investigate those items to determine whether they indicate a misstatement or a control deficiency.⁹ **STANDARD**

There is an operational consequence that firms consistently underestimate. Coverage creates work. A team that moves from a sample of 60 journals to a full population with 900 flags has changed the shape of the engagement, and the budget has to change with it, and somebody has to tell the client that the fee assumption made in October no longer describes the job. My own view is that a firm which deploys population analysis without re-planning the engagement will produce either an unworked population or an unprofitable job, and usually both.

3.4 What to leave alone for now

Two things should stay out of the first wave. Anything that drafts a conclusion, as distinct from drafting a description of work performed, belongs there because a generated conclusion in a working paper is indistinguishable in the file from a reasoned one. Anything that generates evidence, as distinct from summarising it, belongs there for reasons Chapter 4 sets out.

Neither is permanent. In both, the cost of an error is high and the ability to detect it is low, and current guidance offers a firm nothing to stand behind if it goes wrong.



Tax and technical research. The safest deployments are the ones where an authoritative answer already exists to check against.

CHAPTER FOUR

Six ways this fails in an audit

Five of the six failures below leave a file that looks complete. That is what makes them worth naming in advance.

The risks published by the regulator are organised around the tool. The taxonomy below is organised around the file, because the file is what an inspector reads and what a court would read. I am proposing the grouping rather than reporting it. **ECAVEO**

THE FAILURE TAXONOMY

FAILURE	WHAT ACTUALLY HAPPENS	FILE LOOKS COMPLETE
Evidence substitution	A generated output is filed as evidence, with no record of the procedure applied to it	Yes
Completeness illusion	The tool reports accurately on the material supplied, and the population it was given is never tested	Yes
Scepticism decay	The tool offers an explanation, the team tests it without forming an independent expectation	Yes
Documentation without basis	Working papers describe work and conclusions in fluent language that no underlying evidence supports	Yes
Silent tool change	A supplier update alters behaviour mid-engagement, and results before and after are treated as comparable	Yes
Synthetic evidence	Fabricated invoices, confirmations or recordings enter the evidence chain from outside the firm	No, if authenticity is tested

TABLE 3 The right-hand column is the point of the table. A failure that leaves a complete-looking file will not be caught by review, so it has to be prevented by design.

4.1 Evidence substitution is the one an inspector will find

The test an inspector applies is whether an experienced auditor with no previous connection to the engagement could understand the work performed, the evidence obtained and the conclusions reached. A file containing a generated summary, with no record of what was checked, fails that test whatever the quality of the summary.

The remedy is a documentation habit. No technology control will supply it. Where a tool contributed to a working paper, the paper records which tool, for what purpose, what the output was checked against, what proportion was checked, and who performed the check. The Council's June 2025 guidance points in the same direction, describing documentation of the tool's function, its appropriateness for the intended use, its development, its governance and whether its outputs remain appropriately explainable.⁶ **REGULATOR**

4.2 Completeness illusion

Completeness is the assertion these tools are least equipped to support and most likely to appear to support. A system answers about the documents it holds, while the team wants to know whether something exists at all.

The design requirement follows directly. Any tool used on a document population reports what it read, what it could not parse, what it expected to find and did not, and the period it covered. Requiring a denominator converts an unsafe negative into a testable one, and it costs nothing at build time.

DESIGN RESPONSE

Require every extraction tool to report its denominator. A negative finding without a denominator is not a finding.

4.3 Scepticism decay, and why the research is thin

The concern most often raised about these tools in audit is that they narrow inquiry, because a fluent explanation invites testing and discourages independent expectation. There is experimental evidence here, and it points somewhere more interesting than the concern anticipates. A study published in 2022 found that auditors displayed aversion, proposing smaller adjustments to management's estimates when contradictory evidence came from an artificial intelligence system than when the same evidence came from a human specialist, which the authors concluded could impair audit quality.¹⁷ **PEER-REVIEWED**

Read carefully, that finding does not comfort anyone. It says that the presence of a machine in the evidence chain changes how an auditor weighs what the machine produces, in a direction the auditor is unlikely to notice. Whether generative systems produce the same effect, or the opposite one, is an open question, and the field evidence is thinner than the experimental evidence. **EVIDENCE GAP**

The international board appears to share the concern, since its exposure draft adds requirements on professional scepticism aimed explicitly at confirmation bias.⁸ **STANDARD** I would build the control regardless of how the research settles, because the control is cheap. A tool that surfaces contradictory evidence alongside corroborative evidence costs a little more to build and removes one failure mode entirely.

WHERE EACH FAILURE IS CAUGHT

A	B	C	D
Before use	In the output	In the file	After deployment
Certification, intended-use assessment and data-validation requirements stop unsuitable tools reaching an engagement.	Denominators, source pointers, stated uncertainty and contradictory evidence prevent illusion and decay.	A record of what was checked, by whom and against what turns reliance into something an inspector can test.	Version records, regression tests and quality monitoring catch silent change.

FIGURE 2 Each failure has one natural interception point. Firms tend to invest heavily in the first and lightly in the fourth, which is the reverse of what the thematic review found was missing.

Ecaveo framework. **ECAVEO**

4.4 Synthetic evidence, arriving from outside

The sixth failure is the only one a firm does not cause. Documents, confirmations and now live video can be fabricated well enough to deceive a competent professional. The most widely reported case involved an engineering group's Hong Kong office, where an employee was deceived by a video conference using fabricated likenesses of colleagues and made 15 transfers totalling roughly \$25m.¹⁶ **REGULATOR**

That was a payment fraud. It is worth citing as an illustration and no further. The read-across is that a confirmation received by email, a scanned invoice and a recorded management representation are all now cheaper to fabricate than they were, which raises the value of procedures that establish authenticity independently. The board's drafting on authenticity points the same way.⁸ **STANDARD**

CHAPTER FIVE

Certification, and the control model

The regulator named certification as one of four layers of confidence. Its own inspection work found that no firm was monitoring what these tools did to audit quality after deployment, which is the harder gap to close.

Controls should scale with what the output contributes to. How sophisticated the tool is does not enter it. A summary that helps a manager prepare for a meeting and a summary that supports a conclusion on revenue recognition may come from the same system and belong in different tiers.

RISK TIERS AND MINIMUM CONTROLS

TIER	TYPICAL SCOPE	MINIMUM CONTROLS
One	Internal, non-client material, with reversible assistance	Approved tool, training, source checking and human review before anything is issued
Two	Client-confidential information, or internal workflow decisions	Registered use case, data-flow review, evaluation, logging and a named reviewer
Three	Work contributing to a working paper or a client deliverable	Certification for the intended use, documented testing, data-validation requirements and engagement-level records
Four	Work contributing to an audit conclusion, a tax position or a regulated report	A senior accountable owner, independent verification, monitoring after deployment, regression testing and a tested rollback

TABLE 4 Tier is set by what the output contributes to. A tool used at more than one tier is certified for the highest.

5.1 What a certification record should contain

The thematic review describes what was missing and leaves good practice undefined, so the following is my own construction from those findings. **ECAVEO** A certification record should state the tool and version, the intended use and the uses excluded, the data the tool requires and the validation expected of it before use, the testing performed and by whom, the residual limitations, the approver and date, the review date, and the events that trigger recertification.

Two of those deserve emphasis because the review found them weak. Data-validation requirements should be explicit, since some certification templates did not require teams to check the completeness and accuracy of input data before use.⁷ **REGULATOR** Recertification triggers should be defined in advance. Leaving it to a tool owner to raise is how the review described current practice.

5.2 Monitoring after deployment is the missing layer

The review found that no firm performed formal monitoring to quantify the effect of these tools on audit quality, and that the monitoring which did exist was largely restricted to the more widely used tools and driven by licensing.⁷ **REGULATOR** Closing that gap does not require a large programme. It requires a small number of indicators reported on a regular cycle.

Four indicators do most of the work. The proportion of engagement teams using the tool for its certified purpose and not an adjacent one. The rate at which reviewers correct or reject tool output. The results of a standing regression set rerun after every supplier update. The number of engagements where the tool was disabled or abandoned mid-audit, which is the most informative signal a firm can collect and the one least often recorded. **ECAVEO**

5.3 The six technical controls

Underneath certification sit six controls consistent with the security guidelines, the generative risk profile and data protection accountability expectations.^{1,2,3} **STANDARD**

1. Enterprise identity and engagement-level entitlements are enforced before retrieval and before any tool call.
2. Contracts define training use, retention, hosting region, subprocessors, deletion, incident notification and audit rights.
3. Agents operate with least privilege and short-lived credentials, and any consequential action requires explicit approval and can be reversed.
4. Every production system records the model, prompt, retrieval corpus, tools, permissions and evaluation version behind each output.
5. Prompt injection, hostile documents, sensitive-data disclosure and unsafe output handling are tested before scale, using material of the kind the system will meet.
6. Generated work passes the same review and records controls as work produced without it.

ILLUSTRATIVE SCENARIO

The regression set that paid for itself

An invented firm certifies a contract-review tool for revenue recognition work in October and builds a regression set of 40 agreements with known answers. In February the supplier releases an update.

Rerunning the set takes 90 minutes and shows recall unchanged on 38 agreements and materially worse on the two containing variable consideration. The firm restricts the certified use, notifies the three engagements affected and raises the issue with the supplier.

Without the regression set the change stays invisible until an inspection, and by then nobody can say when the behaviour altered. Ninety minutes, four times a year.



PART THREE

Proof, economics and delivery

CHAPTER SIX

Proving it, and the evidence gate

An evaluation designed after people have used the tool will confirm what they already believe. Write the protocol first, and blind the review.

THE SIX DIMENSIONS OF AN EVIDENCE GATE

DIMENSION	WHAT IS MEASURED	THE GATE
Outcome	Cycle time, coverage, search time or throughput	Improves the intended workflow against a real baseline
Quality	Accuracy, completeness, citation support, corrections and rework	No material regression, with critical errors below a stated threshold
Scepticism	Challenge quality, override behaviour, escalation and reviewer understanding	Reviewers form independent expectations and can explain the result
Security and privacy	Permission leakage, client-data exposure, unsafe actions and red-team findings	No unresolved critical finding, with residual risk accepted by a named owner
Economics	Licence, compute, integration, certification, review and support cost	Total cost is justified by verified benefit, including the cost of checking
Operations	Reliability, supplier changes, monitoring, support and rollback	A named owner and tested operational controls are in place

TABLE 5 A pilot passes when every dimension clears its gate. Partial passes are how firms end up scaling something that suits the team that chose it.

6.1 Three pilots worth running

A board minute summariser, scoped to extracting audit-relevant matters with citations and an explicit completeness check, tested on recall of known items, false positives, reviewer time and traceability to source.

A revenue contract reviewer, scoped to identifying clauses relevant to revenue recognition in a controlled evidence set, tested on clause recall, the rate of unsupported conclusions, consistency across reviewers and expert review effort.

A technical accounting assistant, scoped to answering only from approved current standards and firm guidance, tested on citation accuracy, on jurisdiction and effective-date accuracy, and on how often it declines when it should.

Each should be reviewed by an experienced professional who did not run the pilot and does not know which output came from the tool. Blinding costs a little and removes the largest source of bias in an internal evaluation.

6.2 The minimum assurance record

Six things, for every production system. Purpose, intended users, benefit, owner and review date. Allowed and prohibited uses, data classes, jurisdictions and actions. Model, supplier, prompt, retrieval configuration, connectors, permissions and dependencies. Risk assessments, contractual controls and professional decisions taken. Evaluation sample, method, results, limitations and approval. Monitoring, feedback, incidents, changes, rollback and retirement.

6.3 Scale, redesign or stop

Scale where the benefit is achieved, quality is maintained, residual risk has been accepted by a named owner and ownership is funded. Redesign where a value signal exists and quality, review burden, workflow fit or controls miss the gate. Stop where no verified benefit appears, the risk is unacceptable, the fit is weak, or a simpler method performs better.

How many programmes have used the third outcome even once?

CHAPTER SEVEN

The economics of verification

This technology moves cost from producing work to checking it. In audit the checking is performed by the people the firm can least afford to occupy.

A business case that counts hours saved in production and stops there is counting the wrong side of the ledger. Generation becomes nearly free, and the binding constraint moves to verification, which in an audit practice is performed by managers, directors and partners. **ECAVEO**

The test can be applied to any proposal in about a minute. Take the cost of producing the work by hand, the cost of generating it, and the cost of checking the generated version to the standard the work requires. Where generation plus verification comes in below production, the case holds. Where it does not, the firm has bought a faster way to create review burden, and it will feel it at the manager grade.

WHY SOME USE CASES PAY AND OTHERS DO NOT

1 Cheap to check Retrieval with citations, or extraction with pointers to source. A reviewer agrees a sample back and tests the rest.	2 Costly to check An unsourced synthesis. Checking it means performing the work again, which is the work it replaced.	3 Impossible to check A generated conclusion. There is nothing to agree it back to, because the conclusion is the output.
--	--	--

FIGURE 3 Verification cost is a design property. Source pointers move a use case from the second column into the first.

Ecaveo framework. **ECAVEO**

7.1 The cost lines that are missing

Three costs are routinely absent from a pilot business case. Senior review time, described above. Certification and assurance work, meaning testing, red-teaming and supplier diligence, with the records that support them, which the thematic review shows most firms were not yet performing to a consistent standard.⁷ The cost of supplier change, meaning the regression testing and recertification that follows every material update, which recurs indefinitely.

My own rule of thumb, offered as practitioner judgement rather than measured fact, is that a tier four deployment carries ongoing certification, monitoring and review costs of a similar order to its licence cost. **PRACTICE** A firm budgeting only for licences will be surprised twice.

7.2 The pricing question the profession has not settled

Efficiency in an hourly market creates a pricing problem that audit has not resolved and that the adjacent professions are arguing about openly. Audit fee pressure is already present. Public interest entity audit fees in the UK totalled £1.47bn in 2024, up 4 per cent, with the largest four firms earning 91 per cent of those fees while firms outside that group held 40 per cent of the engagements.¹⁸ **REGULATOR**

Audit quality itself is now reported firm by firm, with the 2026 review published in July of that year, so a firm's inspection results and its technology decisions are visible to the same audit committee in the same season.²¹ **REGULATOR** A firm that reduces the cost of an engagement and holds the fee is improving margin, which is legitimate and not indefinitely sustainable in a market with an active regulator and informed audit committees. My own view, which no source states, is that the profession will arrive at disclosure of tool use in the engagement letter well before it arrives at a settled position on pricing, and that firms should prepare for the first conversation.



Close operations and exception approval. The posting stays with a person, whatever produced the recommendation.

CHAPTER EIGHT

Ninety days, then nine months

The first month is observation, inventory and the certification framework. Choosing a vendor first produces a tool nobody is allowed to use.

THE DELIVERY SEQUENCE

PHASE	TIMING	OUTPUTS
Discover	Weeks 1 to 4	Workflows mapped to quality management controls, data boundaries, methodology touchpoints and an inventory of current use
Prove	Weeks 5 to 10	Three pilots on curated evidence sets, with blinded expert review and written stop conditions
Certify	Weeks 11 to 18	Tool approval, documentation, versioning, monitoring indicators and training requirements
Scale	Months 5 to 9	Deployment by service line and engagement type, with quality indicators and inspection readiness

TABLE 6 The Certify phase is the one that distinguishes audit from every other sector in this series, and it is the phase the thematic review indicates firms find hardest.

8.1 The first 30 days

Interview the audit, tax and advisory leadership and observe real work. Inventory what is already in use, including tools embedded in software the firm already licences, and including the unsanctioned use that survey evidence suggests is widespread.¹⁴ Map data boundaries, contractual duties, entitlements and the professionals who actually decide.

Publish interim guidance on safe use within the first fortnight, together with a simple route for proposing an experiment, because a shadow estate grows while a programme designs itself. Select three pilots with baselines, representative evidence sets and written stop conditions. Demonstrate one bounded workflow end to end, so that partners see capability and controls in the same room.

8.2 What success looks like at nine months

Scaled workflows with measured quality or time improvement against the baselines set in month one. A single intake route, certified tools, documented use cases and monitored supplier changes. Professionals who can explain what a tool cannot do, agree its output back to source and escalate a failure. Engagement partners who can show an inspector the basis on which a tool was relied on. A backlog separating proven opportunities from research and from uses the firm has decided against. The government’s own playbook sets out ten principles which, although written for departments, translate cleanly to a partnership and are a reasonable check on a firm’s own policy.⁴ **STANDARD**

PRACTICE
Start the certification framework in week two, not week 11. Pilots run against a framework produce evidence. Pilots run before one produce anecdotes.

CHAPTER NINE

The session to run with the board

Eight questions, applied to any proposed use. A use case that survives all eight is worth certifying.

1. What observable outcome should improve, for whom, and by when?
2. What is the current baseline, and where does the work wait, repeat or fail today?
3. What data is involved, and what client and professional obligations apply to it?
4. What error would matter, who would be harmed, and how would anyone detect it?
5. Who reviews the output, who decides, who can override, and who owns the conclusion?
6. What simpler method should this be compared against?
7. How will a change in the model or the supplier be detected during an engagement?
8. What evidence would justify certifying this, and what evidence would stop it?

The fourth and the eighth do most of the filtering. A proposal whose sponsor cannot describe the error that would matter has not been thought through, and a proposal with no stop condition has already committed the firm.

ONE THING TO DECIDE IN THE ROOM

Name the person accountable for certification. Not the technology, not the supplier relationship, the decision that a given tool may be relied on for a given purpose. The regulator's guidance is addressed principally to central methodology and tooling functions, which is a clear indication of where that person should sit.⁵

CHAPTER TEN

Where the evidence is thin

Four things in this paper would change if the evidence changed, and two of them concern claims the profession currently repeats without support.

10.1 The scepticism question

Experimental work has measured how auditors weigh evidence produced by an artificial intelligence system, and found aversion where the concern predicted deference.¹⁷ **PEER-REVIEWED** What has not been measured is whether generative systems behave the same way, and whether the effect survives outside a laboratory into a live engagement. **EVIDENCE GAP**

A firm could answer the second question for itself. Run a blinded comparison in which one group forms an independent expectation before seeing tool output and another does not, and measure the difference in the questions asked of management. That is a week of work and would produce evidence about this firm, which is the evidence a firm actually needs.

10.2 The quality question

There is evidence that earlier generations of this technology improved audit quality. A study of 36 large audit firms, using employment data and more than 310,000 curricula vitae, reported that a one standard deviation increase in recent artificial intelligence investment was associated with a 5.0 per cent reduction in the likelihood of a restatement.²⁵ **PEER-REVIEWED** That is an association at firm level over a period preceding generative tools, so it tells a partner something about the direction of the technology and very little about the system a supplier is currently selling.

For generative systems specifically I could find no published study establishing an effect on audit quality in either direction, and no enforcement decision in which the misuse of such a tool was the stated ground of sanction. **EVIDENCE GAP** **PRACTICE** My own view is that the inspection cycle will supply the first real evidence, and that it will concern documentation.

10.3 The claims the profession repeats

Two figures circulate widely and neither should be relied on. Investment announcements from the largest firms are self-disclosed commitments. Nobody has audited that spend.¹⁵ **VENDOR RESEARCH** Success rates reported by firms about their own programmes measure expectation. Where a survey does not publish a sampling frame and a response rate, a partner reading it should treat it as sentiment.



The public interest purpose. Everything in this paper exists to protect the moment at which somebody signs.

10.4 The pipeline question

The entry-level position is changing and the direction is contested. In the United States, accounting degrees awarded fell to 55,152 in the 2023 to 2024 academic year, a decline of 6.6 per cent, while candidates sitting the professional examination fell from 42,626 in 2023 to 28,082 in 2024, against which spring 2025 enrolment in accounting programmes rose 12.4 per cent.¹⁹ **SURVEY** Equivalent UK figures I could not verify to a standard worth printing. **EVIDENCE GAP**

The concern most often voiced is that automating first-year work removes the apprenticeship through which judgement is formed. It is a serious concern and it remains an argument. What a firm can do is deliberate. Assign reasoning tasks that a tool cannot perform, review the reasoning itself, and measure whether the capability is forming. That is cheap, and it is testable within one intake.

References

Guidance in this area remains subject to jurisdiction, to the facts of an engagement and to later revision. Where a source is a survey or claim from a firm with a commercial interest in the result, the reference says so.

1. Information Commissioner's Office. *Guidance on AI and data protection*. Page last updated 15 March 2023 and flagged by the Commissioner as under review. <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/guidance-on-ai-and-data-protection/>
2. National Cyber Security Centre. *Guidelines for secure AI system development*. 27 November 2023, co-sealed by agencies from 18 countries. <https://www.ncsc.gov.uk/collection/guidelines-secure-ai-system-development>
3. National Institute of Standards and Technology. *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*, NIST AI 600-1. 26 July 2024. <https://doi.org/10.6028/NIST.AI.600-1>
4. UK Government. *Artificial Intelligence Playbook for the UK Government*. 10 February 2025. <https://www.gov.uk/government/publications/ai-playbook-for-the-uk-government/artificial-intelligence-playbook-for-the-uk-government.html>
5. Financial Reporting Council. *Generative and Agentic AI Guidance*. 30 March 2026, described by the Council as the first such guidance issued by a global audit regulator. <https://www.frc.org.uk/library/standards-codes-policy/audit-assurance-and-ethics/guidance/ai-in-audit/>
6. Financial Reporting Council. *AI in audit: illustrative example and documentation guidance*. 26 June 2025. <https://www.frc.org.uk/library/standards-codes-policy/audit-assurance-and-ethics/guidance/ai-in-audit/>
7. Financial Reporting Council. *Thematic Review: Certification of Automated Tools and Techniques*. June 2025, covering the six largest UK audit firms. https://media.frc.org.uk/documents/Thematic_Review_on_the_Certification_of_Automated_Tools_and_Techniques.pdf
8. International Auditing and Assurance Standards Board. *Exposure Drafts on Audit Evidence and Risk Response*, covering ISA 330, ISA 500 and ISA 520. Approved June 2026, released August 2026, comments close 15 December 2026. <https://www.iaasb.org/consultations-projects/audit-evidence-and-risk-response-isa-330-isa-500-isa-520>
9. Public Company Accounting Oversight Board. *Amendments Related to Aspects of Designing and Performing Audit Procedures that Involve Technology-Assisted Analysis of Information in Electronic Form*. Release 2024-007, adopted 12 June 2024, approved by the SEC on 20 August 2024, effective for fiscal years beginning on or after 15 December 2025. <https://pcaobus.org/oversight/standards/implementation-resources-PCAOB-standards-rules/amendments-related-to-aspects-of-designing-and-performing-audit-procedures-that-involve-technology-assisted-analysis-of-information-in-electronic-form>
10. International Auditing and Assurance Standards Board. *ISA 500 Series project*, covering ISA 501 and ISA 505. Project proposal approved March 2026, exposure draft in development. <https://www.iaasb.org/consultations-projects/isa-500-series>
11. Institute of Chartered Accountants in England and Wales. *Using generative AI*. <https://www.icaew.com/technical/technology/artificial-intelligence/generative-ai-guide/using-generative-ai>
12. Pierpoint, H. *Testing Gen AI platforms: can they do audit work?* ICAEW Audit and Assurance Faculty, September 2025. <https://www.icaew.com/insights/viewpoints-on-the-news/2025/sep-2025/testing-gen-ai-platforms-can-they-do-audit-work>
13. Chartered Accountants Worldwide with Ipsos UK. *Survey of chartered accountants on artificial intelligence*. Published 11 April 2025. Sample of 2,718 across 48 countries and 13 institutes. <https://charteredaccountantsworldwide.com/ai-transformation/>
14. Thomson Reuters Institute. *Future of Professionals Report 2026*. 22 June 2026. Survey of 1,816 professionals across 62 countries, fielded March to April 2026. Vendor-adjacent survey covering several professions, audit among them.
15. Investment announcements by the largest audit and advisory firms, 2023 to 2025. PwC announced \$1bn over three years on 26 April 2023, EY announced \$1.4bn alongside its EY.ai platform in September 2023, and KPMG announced a \$2bn technology and artificial intelligence alliance in July 2023. Self-disclosed commitments, not audited spend.
16. Deepfake fraud at an engineering group's Hong Kong office, disclosed by Hong Kong Police on 2 February 2024. Approximately HK\$200m, or about \$25m, transferred across 15 payments. A payment authorisation fraud rather than an audit evidence fraud.
17. Commerford, B. P., Dennis, S. A., Joe, J. R., and Ulla, J. W. *Man Versus Machine: Complex Estimates and Auditor Reliance on Artificial Intelligence*. *Journal of Accounting Research*, 60(1), 2022, 171 to 201.
18. Financial Reporting Council. *Audit market and competition developments, 2025 update*. December 2025. https://media.frc.org.uk/documents/Audit_market_and_competition_developments_2025_Update.pdf
19. American Institute of Certified Public Accountants. *2025 Trends report*. Published 27 October 2025, reported in the *Journal of Accountancy*.
20. American Institute of Certified Public Accountants. *SAS No. 142, Audit Evidence*. Issued 9 July 2020, effective for periods ending on or after 15 December

2022. Recognises automated tools and techniques as a means of obtaining audit evidence.
21. Financial Reporting Council. *Annual Review of Audit Quality 2026*. 22 July 2026. <https://www.frc.org.uk/news-and-events/news/2026/07/>
 22. International Organization for Standardization and International Electrotechnical Commission. *ISO/IEC 42001:2023, Artificial intelligence management system*. December 2023.
 23. International Auditing and Assurance Standards Board. *ISQM 1, Quality Management for Firms that Perform Audits or Reviews of Financial Statements, or Other Assurance or Related Services Engagements*. Systems required to be designed and implemented by 15 December 2022. Paragraph 32(f) covers technological resources.
 24. International Auditing and Assurance Standards Board. *ISA 220 (Revised), Quality Management for an Audit of Financial Statements*. Effective for periods beginning on or after 15 December 2022. Paragraph 25 covers engagement resources.
 25. Fedyk, A., Hodson, J., Khimich, N., and Fedyk, T. *Is artificial intelligence improving the audit process?* *Review of Accounting Studies*, 27(3), 2022, 938 to 985. Firm-level association covering a period preceding generative tools. <https://link.springer.com/article/10.1007/s11142-022-09697-x>

ECAVEO RESPONSIBLE AI SERIES

Sufficient and Appropriate

Paper two of three. Version 1.0, August 2026.

ALSO IN THIS SERIES

Paper one. The Traceable Thesis. Artificial intelligence across the private equity investment lifecycle.

Paper three. Verified Authority. Artificial intelligence, privilege and the duty to the court in law firms.

THE AUTHOR

Paul Forrest is a fractional Head of AI working with private equity and venture capital portfolio businesses on the deployment of artificial intelligence.